

Money Matters—In a Good Institutional Environment

MONEY MATTERS—BUT NOT IN THE WAY WE thought. Projects financed by foreign aid are often highly visible and important successes—roads and highways, schools and health clinics, irrigation infrastructure, power plants. But success can be assessed at two levels—at the micro or project level, which typically shows high rates of success or at the macro level of economywide growth and poverty reduction, where, as chapter 1 has shown, there is less visible success. If aid financing is fungible, the benefits of an aid-financed project are only loosely connected with the actual benefits of aid financing. Aid's true effect depends on the crucial (but difficult to assess) question: what would have happened in the absence of donor financing?

Most aid goes to projects to assist particular sectors (figure 3.1). Do these targeted resources reach these sectors? Or is foreign aid fungible? Specifically, if donors finance a \$1 million project in education, does the recipient country spend more on education than if the donors had simply provided \$1 million of budget support? The issue is not corruption, nor is it a question of administrative arrangements for aid flows. Instead, a person—or country—given resources “in kind” will rationally reallocate their other expenditures.

If aid is not fungible, evaluating the overall effect of aid is easy: it is simply the collective effect of individual projects. If primary education is important, more projects in this area should lead to overall improvements. If aid is typically fungible, however, managing and assessing the effect of aid is much more complicated. What you see is *not* what you get. Evaluating aid's effect requires not just examining what happened in

the lifetime of a project, but judging what would have happened had there been no aid.

The first section of this chapter explores in detail what is meant by fungibility. We then look at the available evidence on several dimensions of fungibility. If a donor finances an investment for education, for instance, what confidence does it have that the resources:

- Stick with the government rather than finance, say, tax reductions?
- Finance investment rather than expanding consumption spending?
- Finance education rather than some other sector?


Donors cannot be certain that resources actually go to the area favored by the project. So, we examine the implication of fungibility for evaluating projects.

One possibility is that in practice the actual effect of financing tied to projects is to raise government spending exactly as if general budgetary support were given. This implies that, even when donors are financing projects, an assessment of aid's effect as a financial transfer depends on a judgment of overall government spending—whether expenditures are allocated to activities likely to promote development and how effective the spending is. This involves assessing the rationale for government spending. Again, what you see is not necessarily what you get, as the key question is: what would have happened without government interventions?

What does fungibility mean for analyzing the allocation and efficiency of government spending and for managing and evaluating development assistance? In countries with sound policies, appropriate allocations of expenditures, and effective services, donors can provide large amounts of assistance as general budget support, knowing the resources will be well used. In cases where there is agreement about allocation but efficiency is low, aid projects and financing should be evaluated not just as money for a particular project but also for what they contribute to improving the overall efficacy of government expenditures (see chapter 4). Where donors and governments do not agree on the allocation of expenditures and spending is not likely to be effective, the best approach is to reduce funding and increase support for policy dialogue and institution building—until donors are convinced that their funds will contribute to development.

Most aid goes to projects—and most project aid goes to social, economic, and administrative infrastructure.


Figure 3.1 Distribution of Aid, by Type and Sector



Source: OECD 1998.

Fungibility is an issue only if the objectives of donors and recipients are different.

Figure 3.2 Full Fungibility



Source: Hypothetical illustration.

Fungibility?

FIRST, A SHORT BUT INFORMATIVE DETOUR. MANY GOVERNMENT programs transfer income to individuals—income that is tied to specific purchases or that is transferred “in kind” rather than in cash. A prime example is a food stamp program that provides vouchers to be redeemed for the purchase of a limited list of food items. Even if the program is free of corruption (so that every food stamp dollar is redeemed for food), does such a program have a greater effect on people’s food consumption than if they had been given cash?

Compare two views. The naive “what you see is what you get” view is that, if food stamps bought food, food consumption increased dollar for dollar. The “fungibility” view is that food stamp recipients are sophisticated and have their own objectives, in which case the recipient would reduce expenditures on food from their non-food stamp income dollar-for-dollar with increases in food stamp expenditures. So a dollar’s worth of food stamps, even when spent on food, will increase food consumption by exactly the same amount as would have a dollar given in cash.

Which view is right? For food stamps there is solid evidence. Puerto Rico “cashed out” food stamps in 1982, and evidence suggests that this had little or no effect on food expenditures (Moffitt 1989). In the United States two randomized experiments of “cash out” found no effect on food expenditures in Alabama and only a small effect in San Diego (Fraker and others 1995). With food stamps the empirical evidence is overwhelming—that the naive view is wrong and the fungibility view is (almost exactly) right. But how does this relate to foreign aid?

A simple diagram helps illustrate three points about the fungibility of foreign aid (figure 3.2). Suppose that a government spends on only two goods, roads and schools, and in its budget it chooses to spend R on roads and S on schools (point X on the budget line). Now suppose that a donor is willing to make a grant of amount A to finance roads. If donor money is not at all fungible, this entire amount would be spent on roads. Thus the government would spend $R+A$ on roads, and the same amount as before on schools (S). But if the money is completely fungible, the government’s new budget is effectively the pre-grant budget plus A (as long as the government meets the condition of spending more than A on roads). In this example the government would then spend R' on roads (more than the pre-grant R and more than A but less than $R+A$) and S' on schools. This illustrates three points.

First, all of the aid (A), or its equivalent, is spent on roads in an administrative sense in all cases, but the effect of aid spent on roads on the *total* spending on roads could range anywhere from zero to the total amount A .


Second, full fungibility does not mean that none of the aid gets spent on roads, just that the effect on total road spending of providing amount A for roads is exactly the same as if the amount were provided as general budget support. The practical importance of fungibility depends on whether the donor and recipient have the same objectives. If the donor simply wishes to expand the budget available to the government but it is administratively convenient to provide funding for roads, then fungibility is unimportant. If, on the other hand, the donor has different preferences and really wanted the amount A to be spent on roads, then fungibility is a problem.

Third, fungibility requires that the amount of aid provided for roads is smaller than that which the government would have spent on roads from its own resources. Suppose that the government would have chosen to spend at point X on roads and schools and then receives road aid of amount A (figure 3.3). Then actual expenditures will increase to point Y , but had the transfer been untied, the government would have chosen point Z . Thus fungibility is less likely if the amount of aid is large relative to the government's budget (as in many aid-dependent economies) or if the item to which the financing is tied is specific and only a tiny fraction of the government budget (for example, project aid made available for an entirely new type of spending).

Most foreign aid goes to projects in which the donor finances either a single large investment (such as a highway) or, more often, a cluster of related investments (rural roads, urban housing, systems of irrigation canals, water supply for a number of villages, programs of school construction). Yet the fungibility of aid financing has been recognized as an issue from the beginning of large-scale aid.¹ The analytical possibility is not new, but what is new is the evidence on fungibility. Suppose that a donor provides \$1 million to a government for an investment in a specific project in, say, education. The three relevant questions are:

- Does government spending rise by \$1 million?
- Does government development spending rise by \$1 million?
- Does education spending rise by \$1 million?

Figure 3.3 Partial Fungibility



Source: Hypothetical illustration.

There is apparently a greater “flypaper” effect with concessional loans than with grants.


Does Aid Increase Government Spending?

Governments have several options when they receive aid. They can increase spending dollar for dollar. They can reduce taxes dollar for dollar. Or they can leave current spending and taxation unchanged and use the aid to reduce the deficit (equivalent to cutting future taxes)—or any mix of the three.

There is a large body of literature on the effect of transfers between one level of government and another in the same country—say, between the central government and provincial (or state) governments. The evidence is that money from higher-level to local governments tends to pass more into government spending and less into tax relief than a pure fungibility view would suggest. This tendency for nontax resources to “stick” to higher levels of government spending and not be passed on as tax relief is called the “flypaper” effect.


With foreign aid there is a wide range of estimates of the size of the flypaper effect. Two studies based on sizable samples of countries find that \$1 in aid translates into much less than \$1 in government spending. Feyzioglu, Swaroop, and Zhu (1997) looked at 38 countries and found that \$1 in concessional loans (such as those from the International Development Association) leads to 63 cents in additional government spending (figure 3.4). For concessional loans and grants combined, the effect is only 33 cents. Thus there is apparently a greater flypaper effect with concessional loans than with grants. In a sample of 46 countries, Cashel-Cordo and Craig (1990) found that bilateral loans had no effect on government spending.

Studies based on smaller samples and individual countries find stronger flypaper effects. Feyzioglu, Swaroop, and Zhu (1997) use a subsample of 14 countries for which they had detailed data over time about government spending by sector (appendix 3). There was a complete flypaper effect: \$1 of aid translated into roughly \$1 of government spending (95 cents for bilateral loans and \$1.24 for concessional loans). In a case study of Indonesia, Pack and Pack (1990) found that \$1 in aid leads to \$1.50 in additional government spending (figure 3.4). Thus, instead of providing tax relief by substituting for domestic resource mobilization, aid may increase taxation by “crowding in” government spending. This “crowding in” is possible if aid helps government mobilize other resources through improved tax collection or better access to commercial funds—or if aid relieves constraints that were limiting government spending.

Figure 3.4 A Dollar's Worth of Aid and Government Expenditure**Estimates from Large Cross National Samples**

Source: Feyzioglu, Swaroop, and Zhu 1998; Cashel-Cordo and Craig 1990.

In large samples, a dollar's worth of aid leads to significantly less than a dollar's worth of government expenditure . . .

Estimates from Individual Countries and Small Samples

Source: Feyzioglu, Swaroop, and Zhu 1998; Pack and Pack 1990, 1993, 1996.

. . . but in individual countries strong "flypaper" effects are sometimes observed.

It is difficult to find the link between aid and development spending.

This supports the view that aid is largely fungible. Fungibility means that a government can use increased resources as it chooses—to increase spending, fund tax cuts, or reduce the fiscal deficit (reducing future taxes). Different societies will make different choices, so it will be possible to find examples (such as Indonesia) in which government spending rises hand-in-hand with aid. But this relationship is not seen uniformly in large samples of countries, which suggests that although there is some flypaper effect, it is not generally true that \$1 in aid increases government spending by a full \$1. Some goes to tax relief or deficit reduction.

Does Aid Increase Development Spending?


Donor financing and funding for projects is typically not structured to finance government spending in general. Instead, it is aimed at particular projects or specific expenditures, such as investment in infrastructure or social services. In developing country budgets these are usually captured in the distinction between “recurrent” and “development” expenditures.

Many developing countries have weak data on how government spending is split between development expenditures and other spending (administration, subsidies, defense, and so on). Thus it is difficult to find the link between aid and development spending in a large sample of countries. A general idea can be gained, however, by examining the relationship between aid and government consumption spending.

In their sample of 14 countries, Feyzioglu, Swaroop, and Zhu find that \$1 in foreign aid typically results in 29 cents of public investment (figure 3.5). These, remember, were countries with a strong flypaper effect, so aid went dollar-for-dollar to government spending. It seems, though, that aid financed the government in general, not the development expenditures that donors typically target. In this sample 29 cents was the exact amount of a typical dollar of government spending from all sources (aid and non-aid) that goes into investment. Thus an aid dollar had exactly the same effect as one from any other source of government revenue.

The strong effect of aid (especially bilateral aid) on government consumption spending is seen in the same large sample of countries used in chapter 1 (figure 3.6). The strong association of more aid with higher government consumption spending (even after controlling for other determinants of government spending) suggests both a flypaper effect (that resources stuck) and fungibility (that allows aid provided for investment to finance consumption spending).


**Figure 3.5 A Dollar's Worth of Aid and Government Investment
("Development") Expenditures**



Source: Feyzioglu, Swaroop, and Zhu 1998; Pack and Pack 1990, 1993, 1996.

The typical aid dollar finances 29 cents of public investment . . .

Figure 3.6 Bilateral Aid and Government Consumption



Source: Calculated from the Burnside-Dollar data set.

. . . most of the rest goes to government consumption.

Empirical studies of Indonesia and Sri Lanka also find that aid has a smaller effect on development spending than on total government spending, so that not all aid passes into investment. Again, the differences across countries are large. In Sri Lanka an additional aid dollar raised government spending by 52 cents. Of that, 38 cents was devoted to development spending and 14 cents to consumption spending. In Indonesia a dollar's worth of aid raised spending by \$1.58, of which development expenditures accounted for 89 cents.


The evidence suggests that, even where aid goes mostly to government spending, the spending funded by aid is largely fungible between consumption and investment, regardless of whether the aid is administratively targeted to specific investment projects.

Does Project Aid Finance Particular Sectors?

Researchers have also tried to address sectoral fungibility—that is, does higher foreign assistance for a particular sector (for example, agriculture or education) raise spending for that sector? There are two ways to answer

A dollar's worth of aid to agriculture leads to less than a dollar's spending on agriculture.

Figure 3.7 A Dollar's Worth of Aid to Agriculture and Spending on Agriculture



Source: Feyzioglu, Swaroop, and Zhu 1998; Pack and Pack 1990, 1993, 1996.

that question: by comparing spending over time within a country and comparing spending across countries. Whatever the method, the answer is the same: it varies. In some countries and sectors aid appears to be completely fungible across sectors, while in other countries and sectors the money seems to “stick.”

Pack and Pack (1990) show that for every dollar increase in donor financing for Indonesian agriculture, expenditures in that sector increase by 92 cents. In contrast, a cross-sectional study of 14 countries finds that agricultural expenditures *decrease* by five cents for every \$1 in aid given to agriculture (figure 3.7).

The net effect of an aid dollar on spending in a specific sector depends not only on the composition of the aid received, but also on how government responds to aid flows. Take education and health, and the results for a cross-sectional sample of 14 countries. The direct effect on education spending of a dollar in aid depends on the allocation of aid among sectors: 8.7 cents of every aid dollar are allocated to education and health. With sectoral fungibility, however, that is not the whole picture, which depends on whether aid for education leads to a reduction in what the government otherwise would have spent on school programs, and whether aid for other sectors causes the government to spend more on education. In the estimates from the 14 countries, the net effect of a dollar's worth of aid is to reduce education and health expenditures by 6.5 cents, even though on average 8.7 cents was devoted to education and health (figure 3.8).


The effect varies enormously by country. In Sri Lanka 5.9 cents of an aid dollar go to education and health, but the estimated net effect is to *reduce* spending by 1.9 cents. Indonesia receives a similar proportion devoted to education and health, 7.2 cents, but the effect is to *increase* spending by 18.9 cents; clearly money earmarked for other sectors was reallocated to education and health. The important point is that government commitment to particular sectors is more important than targeting aid.

These statistics do not tell the whole story of consequences. It may be that the observed changes in public expenditures are exactly what donors wanted anyway. The donor objective, for example, may not be to increase total education spending, but to increase spending on primary education while reducing spending on higher education. But a shift in the allocation of spending requires agreement between donors and the government based on analysis and dialogue.

Government commitment to particular sectors is more important than targeting aid.

A dollar's worth of aid to education may lead to little (or no) additional spending on education.

Figure 3.8 A Dollar's Worth of Aid to Health and Education and Spending on Health and Education



Source: Feyzioglu, Swaroop, and Zhu 1998; Pack and Pack 1990, 1993, 1996.

Budgetary Mechanics of Fungibility

The magnitude of fungibility depends on a country's budgetary structures, the degree to which governments are able to manage their finances, and the extent of donor involvement.

In India's complicated federal fiscal system of transfers between the central government and individual states, for instance, the budgetary system was designed so that the donor financing coming through the central government to projects in specific states would not affect either the projects chosen or the total transfers received by the states. This made fungibility explicit. The Indian budgetary system was designed so that projects were chosen to be part of the public sector investment plan or not, and then financing was obtained for those projects. The evidence from an empirical study suggests that aid to the Indian central government was fully fungible, both between development and nondevelopment spending and across sectors (Jha and Swaroop 1997). Other countries, such as Mexico, have incorporated mechanisms into their budgetary system to create fungibility: the goal of their system is

to ensure that budget allocations are not to be affected by the availability of finance.

The Indian system also determined the amount of investment financing from the central government to each state. In the past, if states received a larger amount of support from donors, their transfer for development spending on capital account from the central government was correspondingly reduced (Jha and Swaroop 1997).

One major weakness of aid in Africa is the lack of strong budgetary control by either government or society over public investment. A case study in Ghana, Malawi, and Uganda revealed that controls were too weak to accomplish any of the three main objectives of budgetary systems: developing aggregate budget discipline, imposing strategic priorities, and encouraging technical efficiency (Campos and Pradhan 1996). The big problem is that budgets are often ad hoc—that is, realized expenditures differ enormously from budgets, and there is a disconnect between the recurrent budget and the development budget. Moreover, donor financing, particularly through projects, is a large component of investment and total expenditures. For instance, donor project financing accounted for 71 percent of capital expenditures in Ghana and 87 percent in Uganda (table 3.1).

Weak systems of budgetary planning, control, and management, combined with large donor contributions, often cause “the bulk of the Public Investment Program to be determined by donor-driven priorities, thereby fragmenting and undermining the priority setting process within the countries” (Campos and Pradhan 1996, p. 29). Lack of government control over project selection and the budgeting process was one of the key problems in African aid identified in a recent study by the Overseas Development Council and implemented jointly with African researchers (van de Walle and Johnston 1996). Their case studies of seven countries found that donors often drove project selection and most projects had little recipient input. For example, in Senegal only 170 of 316 projects in

Realized expenditures differ enormously from budgets.

Table 3.1 Donor Financing of Investment Budgets in Two African Countries (percent)

	<i>Ghana</i>	<i>Uganda</i>
Project financing/capital expenditures	71	87
Project financing/total expenditures	27	43
All donor financing/total expenditures	32	67
<i>Source:</i> Campos and Pradhan 1996.		

the public investment plan had been assessed using the appropriate internal procedures. Donor control likely meant that aid would be less fungible, but it raised other problems—such as a lack of government ownership and commitments to recurrent finance for projects.

Project Assessment with Fungible Money

These differences in fungibility across countries mean that it is pretty certain that what you see is *not* what you get. Moreover, it is difficult to know either before or after donors provide finance the extent of fungibility, and so exactly what was financed. But if money is fungible, evaluating donor-financed projects is beside the point, as their effect is not the same as that of donor financing. To take a hypothetical example: suppose there were 11 discrete projects, each of which cost \$10 million. Suppose, also, that the economic returns on them ran from 100 percent (for project A), to 90 percent (project B), and so on down to 0 percent (project K). Suppose, finally, the country will invest \$50 million from its own resources, while a donor provides \$10 million—in total, enough to finance six projects. There are several possibilities (table 3.2).

Table 3.2 Evaluating the Effect of Projects with Fungibility

	<i>Projects undertaken</i>	<i>Apparent effect of aid-financed project</i>		<i>Actual effect of aid financing</i>	
		<i>Project</i>	<i>Return</i>	<i>Project</i>	<i>Return</i>
Base case	ABCDE				
Donor finances A, the highest-return project, freed resources finance F, the marginal project	ABCDE, F	A	100 percent	F	40 percent
Donor finances marginal project chosen by the government	ABCDE, F	F	40 percent	F	40 percent
Donor finances A, the highest-return project, and freed resources finance K, a low-return project	ABCDE, K	A	100 percent	K	0 percent
Donor finances A, the highest-return project, and freed resources finance consumption spending, or reduce taxes	ABCDE	A	100 percent	None	Unknown, could be large or small

Note: Based on an entirely hypothetical example in which there is a set of possible projects (each costing \$10 million) with different economic rates of return—and in the absence of aid financing, the government chooses to implement five projects.

Source: Hypothetical illustration.

First, the government agrees that the donor will finance the highest-return project (A). Thus the apparent rate of return is 100 percent: that is, when the donor assesses the benefits it will find that its aid has been hugely successful. What really matters in evaluating the effect of aid, however, is what the government does with the resources it otherwise would have devoted to the project. It could do the next best project not yet financed, which would be F (with a return of 40 percent), since it is already committed to B, C, D, and E. Or the government could choose to do a white elephant project (K) with political payoffs but with a return to the economy of zero. It could also choose to reallocate the resources from investment to consumption. What is the return on additional consumption expenditures, high or low? Often consumption spending, such as that on recurrent inputs or operations and maintenance, is not financed by donors but has enormously high returns in keeping existing investments productive. So a shift from investment to consumption could have a huge return, say, 200 percent. Alternatively, the consumption could be the conspicuous consumption of the elite and have zero return. Or the aid could have financed tax relief, which as a “project” may have very high returns if it reduces distortionary taxes.

Delivering aid through projects has several advantages for donors. There are administrative advantages to organizing tasks into discrete, manageable units. Investment projects fit well with the notion that increasing investment is the key to development. And by financing particular items (family planning, education, and so on), projects appeal to single-interest coalitions in donor countries supporting foreign aid.

But while the project approach has some advantages for donors, it does not ensure that aid is put to good use. Nor does it facilitate the evaluation of aid effectiveness. If, say, aid finances a road project with a 20 percent rate of return, this information tells us next to nothing about aid effectiveness. Does fungibility imply that projects are irrelevant and should be abandoned as a way of organizing and delivering aid, in favor of broader financing mechanisms? Not at all. Project financing can be used to convey other value added. In many cases the project’s value added was never intended to be the investment itself. Donors often are not trying to influence project selection or directly alter the sectoral composition of spending but are using projects as a vehicle for technology transfer or to build institutional capacity or reform sectoral policies.

Taking the road example, donor involvement may mean that those institutions responsible for roads are improved—or road policies are

While the project approach has some advantages for donors, it does not ensure that aid is put to good use.

Donors are, more or less, financing whatever the government chooses to do.

changed for the better. Projects are a cauldron of money and ideas. Since the money does not necessarily stick with the sector, the design and evaluation of projects should focus on the contribution of ideas to altering sector institutions and policies (see chapter 4).

Public Spending—Quality not Quantity

THE SAFEST ASSUMPTION FOR DONORS IS THAT THEY ARE, MORE or less, financing whatever the government chooses to do. Since donors are financing government in general, they need to consider all public expenditures, in terms of both allocation and efficiency, when deciding the appropriate level and method of aid. So donors must each decide, what makes for an effective public spending program?

Equity and Efficiency

Governments ought to be doing things where there is a convincing rationale for public involvement. Twenty or so years ago, donors willingly financed almost anything in which the government chose to try its hand—textile plants, shoe factories, steel mills, and all sorts of manufacturing. Not only were developing world parastatals financed through donor credits and loans; many government corporations were created because donor financing was available. There has been a clear shift in thinking about the value of government investments in shoe making or textile weaving. Today few people would dispute that these activities can just as easily (and with greater effect) be undertaken by the private sector. In general, the rationale for government intervention is to improve equity and efficiency. But government action that achieves these improvements proves difficult in practice.

Equity. Many government expenditures—even those traditionally justified by their improvements in equity—are not well targeted toward poverty alleviation. For instance, an assessment in Brazil found that social insurance payments in 1980 reached only better-off formal sector workers, so the richest 20 percent of workers received 12 times the benefits that went to the poorest 20 percent. In education, too, equity implications differ widely between primary and tertiary education, for instance. A study in Indonesia found that the richest fifth of households received

only 60 percent of the benefits per person from public spending on primary education as the poorest fifth. Since children in richer households were more likely to remain in school, however, the higher the level of schooling, the greater the benefit to richer households. By senior secondary school, the richest fifth received 17 times the benefits of the poorest fifth (van de Walle 1995).

Efficiency. Justifying government actions to increase efficiency is also not always simple. The existence of market failures must be taken into account. But so, too, must their magnitude, particularly when compared with the likely “government failure” in implementing a proposed remedy. The notion that since market failures were pervasive, they were complete and therefore demanded government action, has been replaced by a nuanced view—that both market failures and government failures are pervasive, but not complete. So, although working from common principles, a country-by-country assessment is needed to determine the magnitude of market failure and the depth of government capacity (World Bank 1997b).

Nevertheless, the principle of efficiency often leads to clear recommendations. For instance, governments, sometimes supported by donor financing, have devoted substantial resources to manufacturing and trade that markets could have provided equally well and for which there was no underlying efficiency rationale. Moreover, governments have often used broad-based subsidies that lack rationale in either equity or efficiency—such as large-scale subsidies to agricultural products or fuels used mainly by richer consumers, which impose an explicit or implicit tax on poorer agricultural producers.

Quality and Quantity


The quality of government expenditures is typically more important than the quantity. While there is evidence about what kinds of public sector outputs are crucial to growth and poverty reduction (basic education and infrastructure, for instance), there is often little or no evidence that increasing public spending would have an impact on these objectives. Just as the effect of aid differs across countries, so too does the effect of increased public spending on improving outcomes. Cross-national and sectoral evidence, as well as individual cases, suggests that inputs like money are often not the most important issue for public sector outputs.

The empirical growth literature finds that government consumption has no positive effect on growth. Some studies find a negative effect, others

The quality of government expenditures is typically more important than the quantity.

The quantity of public spending on health bears little relationship to lower mortality.

Figure 3.9 Public Spending on Health as a Percentage of GDP and Health Outcomes



Note: The number in parenthesis is the residual from a regression of the natural log of under-5 mortality on GDP per capita and other socio-economic determinants of mortality.

Source: Filmer and Pritchett 1997.

zero. The finding that aid largely finances government consumption, which has no positive effect on growth, again helps explain why aid is not fueling growth in many developing countries.

One study used data on economic growth and the allocation of government budgets for 43 developing countries to show that there was no relationship between economic growth and public capital spending (Devarajan, Swaroop, and Zou 1996). The authors found that increasing by one percent the share of public spending devoted to capital (versus current) spending would reduce economic growth by half a percentage point each year. Another study, using a specially created dataset on public investments, found that the share of GDP devoted to public investment had no association with economic growth, after controlling for a standard set of determinants of growth (Easterly and Rebelo 1993). So while one intended purpose of aid might be to increase public investment, the quantity of investment alone bears little relationship to economic growth. Examining the effect of the magnitude of public investment across countries shows that differences in the efficiency with which governments invest is as important in explaining growth as the amount they spend on investment (Pritchett 1996).

The data are similarly agnostic about the relationship between sectoral expenditures and economic growth. Devarajan, Swaroop, and Zou find no association between growth and the share of public spending devoted to education or health. Moreover, they find that both education and health spending per capita are consistently negatively related to growth. Easterly and Rebelo also find that the share of GDP going to public investment in education and health is unrelated to economic growth. There is little evidence from cross-country comparisons that more spending on public investment, or even reallocation across sectors, would automatically translate into better growth outcomes.

This finding is consistent with estimates of public spending efficiency for three major sectors: health, education, and infrastructure. Comparisons across countries suggest that public spending has some (small) effect on health. Ranking countries by the difference between actual child mortality and the level expected for their economic and social determinants of health, the 10 countries with the best performance in health typically spend 1.3 percent of GDP on health—and child mortality is 40 percent lower than would be expected (figure 3.9). The worst 10 have child mortality roughly 40 percent higher than expected, but actually devote a larger share of GDP to public spending on health than the 10 best (Filmer and Pritchett 1997).

In explaining any measure of mortality (infant, child) or life expectancy, regression analysis finds little or no correlation between public spending and health status (Musgrove 1996).

Likewise, in education there is little evidence that higher spending means higher quality schooling. Using assessments of educational performance in math and science based on internationally comparable examinations, one study found that both total and recurrent spending on education as a fraction of GDP was negatively related to test score achievements (Hanushek and Kim 1995). While most agree that a rapid expansion of skills was key to the pre-1997 success of the fast-growing East Asian economies, these countries are not distinguished by the extent of the resources devoted to education during its takeoff and after (public spending on education accounted for about 3 percent of GDP in the Republic of Korea) or by physical inputs such as the pupil-teacher ratio, which in 1985 was 38 in Korea, 31 in the Philippines, and 32 in Sri Lanka (Tan and Mingat 1992).

In infrastructure there are similar differences in cost and performance across countries (World Bank 1994). While the flow of infrastructure services is important for economic growth, it is harder to show that infrastructure expenditures (or even stocks) are important (Hulten 1996).

It will come as no surprise to anyone with experience in developing countries that comparisons show little connection between spending and outcomes. This does not mean, however, that spending cannot influence outcomes. The actual effect of spending in almost every area of endeavor—from infrastructure to education—has been much less than the potential.

When the *what* and the *how well* come together and governments do the *right* things *well*, the results can be amazing. An obvious example is the educational accomplishments in East Asia, where countries such as Korea have gone from widespread illiteracy to educational quality that rivals (and surpasses) that of much richer countries. There are similar examples in health, where the mobilization of public health efforts to attack problems of communicable diseases, diseases carried by pests, and water-borne diseases led to a revolution in health status. Sri Lanka and the state of Kerala in India have life expectancies far exceeding that predicted by their income and education. When the state of Ceara in Brazil was able to mobilize and energize government workers in the health sector, it achieved enormous reductions in infant mortality in a few years.

When the emphasis is on what government should do, and particularly how much it should spend, aid becomes a means to support spending, as

The actual effect of spending in almost every area of endeavor—from infrastructure to education—has been much less than the potential.

Donors must form an opinion on the allocation and effectiveness of a country's public spending.

if spending were the end in itself. Thus aid flows independently of how effectively spending promotes poverty reduction. In many cases, even when donor resources increased investment in the right areas, it did not promote development. Thus it is not just what governments do, but how they do it and how well.

What Should Donors Do?

WHAT ARE THE IMPLICATIONS OF FUNGIBILITY AND OF DIFFERENCES in the quality of public spending for the way donors do business? Many discussions of aid policy concentrate on what donors ought to be financing: girls' education or rural feeder roads, agricultural extension or family planning clinics, power plants or vaccination campaigns. But if money is fungible among sectors (and it often is), then from the narrow view of expanding expenditures it simply does not matter what donors finance. Most development practitioners recognize fungibility, however, and have moved to a broader concept of the effect of donor financing: it is more than just an attempt to expand capital. This raises the question of how donors allocate effort and available resources to best assist each country. As with country-level policies, we return to two themes in choosing the magnitude and mode of assistance: timing and selectivity, and the role of money and ideas.

Form a View . . .

To pursue an effective development strategy, donors must form an opinion on the allocation and effectiveness of a country's public spending. Donors can help only if they know what a country's needs are. In the past decade public expenditure analysis has emerged as a useful tool to assess the value of government spending in different areas. Reviews of public expenditures have focused on how well budgets meet the three criteria of budgeting: fiscal discipline, priorities across sectors, and technical efficiency.

Donors need to be able to form a judgment on how well government spending is working along each of the three criteria. These reviews led to the elimination of many "white elephants." Scrutinizing (and often

cutting) large planned investments that would later generate budget commitments was one fruitful result of public expenditure analysis.

Public spending analysis has also brought to the fore attention on issues within sectors—across subsectors and across functional categories. When countries are faced with fiscal adjustment, two common tendencies are across-the-board cuts in budget items and cuts in everything but payroll—approaches that are inefficient. Analysis has also been effective in bringing into focus subsectoral priorities—say, preserving primary education, less so tertiary, from budget cuts.

The first generation of public expenditure analyses did not go far enough. The analyses generally did not use a framework for spending that established justification for public expenditures or broad guidance about the appropriate composition of spending across sectors (Swaroop 1997). A review of 113 World Bank public expenditure reviews found that few even looked at the role of the state and public-private boundaries (McCarthy 1995). Moreover, very few reviews went beyond expenditures and linked them to specific outputs.

One painful lesson of experience is that government and community ownership of projects is crucial.

... and Use the View

Once donors have formed a view on the allocation and efficiency of public spending, what then? One common donor reaction to the fungibility of money is to try to avoid it by imposing additional conditions on aid financing beyond one project to ensure the additionality of resources at the sectoral level. Bad idea.

One way to ensure additionality is for donors to finance only items or projects that governments are not financing. Common sense tells us, however, that this reasoning is flawed. It would mean that donors finance only what countries and governments do not want to do. And government would have little or no commitment to implementing or, much less, maintaining projects once donors depart. Some governments (purposely or by default) have given donors free rein to do whatever they feel like. Money for such projects would almost certainly be “additional,” but with equal certainty the projects themselves would be pointless, except in those rare cases that the project itself creates commitment. Far better that both recipients and donors be strongly committed to the sustainability of projects. But this is precisely when money will be fungible.

One painful lesson of experience is that government and community ownership of projects is crucial. Trying to make financing additional,

Donors should take it for granted that their financing is fungible because that is reality.

however, implies that there is a disagreement between donors and government about the allocation of expenditures but that donors have enough clout to ensure that governments will follow their recommendations on the overall allocation of spending. This is practically impossible. It is not just a question of saying that if a donor provides an additional \$1 million that spending on the activity financed must increase by \$1 million because spending might have increased by that amount in the absence of the extra aid. To ensure additionality, donors would have to make governments promise that they are spending \$1 million more than they would have had they not received the aid. But they would say that in any case, wouldn't they?

In fact, donors should take it for granted that their financing is fungible because that is reality. This, in turn, requires rethinking how aid is delivered. Efficient delivery in different countries depends on the nature of public spending and the efficiency of the public sector. One measure of public sector efficiency is the index of economic management. And an earlier recommendation was that of two equally poor countries, well-managed countries should receive more finance than poorly managed ones (see chapter 1).

In a country with sound public expenditure management, a larger portion of aid can be in the form of general budget support. This recognizes the reality of fungibility and economizes on the administrative costs of aid. These are important—a typical World Bank project runs about \$1 million in administrative costs. If a project is merely replicating what the country already does relatively well (building another power plant in a well-managed system), there is little value added from administrative costs. Providing general budget support would free these staff and donor resources for other, higher-value added activities.

Problems with public expenditure management are usually one of two types: inefficiency in allocation of resources across activities or inefficiency in the use of allocated resources.

Composition of Spending. Consider first a country that does things reasonably well, but continues to spend on many things that are inefficient because the private market could do these things just as well or the social returns are simply too low. Suppose that a country insists on putting large amounts of public finance into higher education. Donors traditionally want to fund primary education to expand its relative share of the budget. But with fungibility this approach is not rational. Possibly a better approach is “time slice” financing of the education budget.

Developed specifically for this purpose are such instruments as the Sector Investment Program. Donors and recipient governments reach a broad agreement on the goals for the sector—and develop projects and programs to achieve these goals. Donors then finance particular projects within the agreed program or provide general support to the sector budget. Periodic reviews of the Sector Investment Program provide a forum where donors and recipient governments can frankly exchange views on the overall sector spending program. Donors can use this policy dialogue to improve the allocation of expenditures. Another mechanism some have tried is a Medium-term Expenditure Framework that tries to produce a strategic direction for government. Case studies of Ghana and Uganda, however, have found these less successful than hoped. In Ghana, some expenditures are designated as core. Pressures to include more items in the core led to the creation of a “supercore,” which was just as pressured. Both concepts were abolished in 1990. One cannot expect miracles, however, and some early experiences with these innovations are showing natural implementation difficulties.

While it is useful to have an open dialogue, it is unrealistic to expect complete consensus on the important activities for government to undertake. There are still wide divisions of opinion about the right role for government between and within donor and recipient countries. Some people believe that governments should subsidize tertiary schooling. Others do not. Some believe that governments should play a strong role in allocating credit. No, say others. The government should provide electricity, say some. Many disagree. These are not just disagreements about the extent and nature of market failure but about values, like the extent to which governments should take redistributive action. Technocratic advice will not be sufficient to resolve most disputes about allocations across sectors or within sectors across categories of spending. In 1990 Sweden spent three times Japan’s 11.5 percent of GDP on social transfers (Lindert 1994). Given the enormous differences among donor countries in the allocation of spending and of private and public responsibilities, it is hard to believe that a sharp consensus will emerge from the democratic process even among donors, much less between donors and recipient governments.

Effectiveness of Spending. Suppose now that an aid recipient has spending allocations that donors broadly support, but is inefficient in delivering services. If a country has carried out macroeconomic reforms, it should be getting a significant amount of aid. Macroeconomic reform is not technically difficult, but developing an efficient public sector is more challenging.

There are still wide divisions of opinion about the right role for government between and within donor and recipient countries.

In this case donors should largely stick with the project approach. The main objective of projects now is not to increase funding for a particular activity, but to help change the institutions and policies at the micro level that affect service delivery. An assessment of public sector efficiency is needed to make an informed decision about how to deliver aid. Once countries have carried out macroeconomic reforms, those with efficient public sectors can receive budget support, while those with inefficient public sectors would get relatively less money and more ideas.

Where a country has both a severely distorted allocation of expenditures and an inefficient public sector, institutions and policies will be poor and finance will be of little value. The traditional donor response has, nevertheless, been to try to find something useful to finance. Fungibility helps explain why large amounts of aid have had no lasting effect in these highly distorted environments. There is no alternative to patience. But there are a variety of ways that donors can, without large flows, support the creation of an environment for productive public expenditures, the subject of chapter 4.

Note

1. The World Bank's first Chief Economist, Paul Rodenstein-Rodan, emphasized in the 1950s that even when financing was tied to projects, money often remained fungible. Assessments of aid as early as the 1960s raised the issue of fungibility (Little and Clifford 1965).